arXiv:1403.7676v1 [stat.CO] 29 Mar 2014

COMPUTATION OF MAXIMUM LIKELIHOOD ESTIMATES FOR MULTIRESPONSE GENERALIZED LINEAR MIXED MODELS WITH NON-NESTED, CORRELATED RANDOM EFFECTS

ANDREW T. KARL

Adsurgo LLC, Denver, CO

YAN YANG

Arizona State University, Tempe, AZ

SHARON L. LOHR

Westat, Rockville, MD

Abstract. Estimation of generalized linear mixed models (GLMMs) with non-nested random effects structures requires approximation of high-dimensional integrals. Many existing methods are tailored to the low-dimensional integrals produced by nested designs. We explore the modifications that are required in order to adapt an EM algorithm with first-order and fully exponential Laplace approximations to a non-nested, multiple response model. The equations in the estimation routine are expressed as functions of the first four derivatives of the conditional likelihood of an arbitrary GLMM, providing a template for future applications. We apply the method to a joint Poisson-binary model for ranking sporting teams, and discuss the estimation of a correlated random effects model designed to evaluate the sensitivity of value-added models for teacher evaluation to assumptions about the missing data process. Source code in R is provided in the online supplementary material.

NOTICE

This is the author's version of a work that was accepted for publication in *Computational Statistics & Data Analysis*. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in *Computational Statistics & Data Analysis*, [VOL73, May 2014] DOI:10.1016/j.csda.2013.11.019

1. Introduction

Generalized linear mixed models (GLMMs) are popular for fitting data sets with correlated responses from a member of the exponential family, where random effects are used to account for within-subject correlation. They have been used in applications ranging from ecology and evolution (Bolker et al., 2009) to epidemiology (Brumback et al., 2013) to education (Broatch and Lohr, 2012). Computing maximum likelihood (ML) estimates in GLMMs is more complicated than comparable computations in linear mixed models because the log-likelihood function for the parameters in a GLMM contains integrals over the random effects that must be computed numerically. The dimension of the integration depends on the structure of the random effects. Many of the computational methods in current use for GLMMs concentrate on nested generalized linear mixed models, in which the units at the top level of the hierarchy are independent. The nested structure considerably simplifies the computation of ML estimates because it reduces the dimensions of the integrals that need to be calculated and simplifies the inversion of relevant covariance matrices.

Many longitudinal studies of subjects in social sciences or epidemiology, however, involve a non-nested structure of the random effects. In education, students proceed through different classrooms in their academic

career. Thus, the response of a student at time t may depend on the full set of classes the student has taken up to that time. This leads to a multiple membership structure (Browne et al., 2001) in which each student is associated with multiple random classroom effects in sequence. The correlation structure becomes complex in multiple membership models because one student's response is correlated with responses of all the students with whom the student shares classes. Each student may have multiple and mixed responses at each time: for example one response might be a score on a standardized test, while another response might be whether the student is required to attend summer school. Furthermore, allowing binary responses allows for the modeling of a missing data mechanism: one of the responses is whether the student is present for the exam; if the student is absent, then his test score is missing. Jointly modeling the student test scores and missing data mechanism allows for a sensitivity analysis to assumptions about the dropout process (Karl et al., 2013a).

Multiple membership models have also been used in longitudinal studies of health outcomes. Chandola et al. (2005) use them to examine the effects of neighborhood and family on health outcomes, when individuals can move from one household to another and from one neighborhood to another over time. Gueorguieva and Sanacora (2006) discuss the advantages of jointly modeling continuous and ordinal measures of depression over time. Accounting for patients seeing different therapists over time would lead to a non-nested structure for the random effects. In a different type of application, Page and Page (2010) study the relationship between the home field advantage and effects of different referees in soccer. The home field advantage of a team for a specific referee is modeled by a random effect, and it would be desirable to also include the home and away teams as random effects. Typically, the number of goals scored in soccer games is low, and a Poisson distribution may be better for fitting the response than a normal distribution.

All of these applications share the following structure. At each time period, there is a potentially multivariate outcome for each subject where at least one of the responses has a non-normal distribution. Each subject is associated with multiple higher-level units, or each observation is a function of multiple levels of the same random factor: a student belongs to multiple teachers or classrooms, a person belongs to different neighborhoods over time, a patient sees different therapists, a referee officiates for different teams and at different locations, and game outcomes depend on combinations of different teams' effects. The multiple membership pattern induces a complex correlation structure in which each observation of a subject is potentially correlated with those of every other subject sharing any of the higher-level units. In a GLMM, the dimensionality of the integration is related to the complexity of the random effects design matrix. When this matrix cannot be factored across subjects, the integral in the log-likelihood will have the same dimension as the number of random effects in the model. In addition, the random effects associated with different responses for the same subject are correlated. The parameters of interest are the fixed regression coefficients, the empirical best linear unbiased predictors (EBLUPs), the variance components for low- and high-level units, and the correlations among the variance components. Pinheiro and Chao (2006) note the need for an efficient estimation routine for non-nested models such as these.

In this paper, we propose a general method based on the Expectation-Maximization (EM) algorithm for computing ML estimates in multiresponse GLMMs with correlated, potentially non-nested random effects. We exploit the sparse structure of the random effects design matrices while using first-order and fully exponential Laplace approximations for the multidimensional integrals to furnish a balance of accuracy and speed, providing an improvement over the widely used pseudo-likelihood model approximation. Furthermore, we express the estimation routine in terms of the first four derivatives of the conditional likelihood for each response, facilitating its application to new problems. Estimation routines for GLMMs tend to be more application-dependent than routines available for linear mixed models; this paper presents a template for jointly estimating multiple GLMMs without assuming any special structure for the random effects design matrices corresponding to each response. R code to fit a non-nested Poisson-binary model and a multiple membership binary response is presented in the online supplement, along with sample data. The code provides a foundation for readers wishing to apply the routine in new settings, since the challenging programming aspects have already been addressed.

In Section 2, we describe the multivariate generalized linear mixed model treated in this paper, illustrate the model with a Poisson-binary example, and summarize previous work on computational methods. Section 3 presents the estimation of the joint model via an EM algorithm. Section 4 discusses the techniques that are necessary in order to compute the EM estimates for large data sets. Applications to sports rankings and student test scores are given in Section 5. A simulation study to evaluate the performance of various approximations appears in Section 6. Section 7 discusses the results and outlines possible extensions.

2. Multiresponse Generalized Linear Mixed Model

In this section we describe the model adopted for the potentially non-nested models studied in this paper. We adopt a correlated random effects model (Fieuws and Verbeke, 2006; Lin et al., 2009; Karl et al., 2013a) in which the multiple observations on a subject are assumed to be independent after conditioning on distinct, but correlated, random effects. This is an extension of a shared parameter model (Wu and Carroll, 1988; Rizopoulos et al., 2009) in which the different responses of a subject share the same random effects.

2.1. **The Joint Model.** Suppose there are n subjects, and between 1 and q response variables are measured on each subject. The data may be unbalanced: it is possible that subjects may have repeated or missing measurements on any one of the q response variables. Let $n_{(i)}$ be the number of observations on response variable i, and let $\mathbf{y}_{(i)}$ denote the vector of those measurements. Response i may be either continuous or discrete, and different responses may be of different types. If only one response is measured, $q = 1$ and the vector $\mathbf{y}_{(1)}$ consists of the data for all subjects. Let $\mathbf{b}_{(1)}, ..., \mathbf{b}_{(q)}$ be distinct yet possibly correlated vectors of random effects which are concatenated (and possibly permuted) into a single vector, $\mathbf{b}$, and assume that $\mathbf{b} \sim N(\mathbf{0}, \mathbf{G})$. The random effects $\mathbf{b}_{(i)}$ may be used to model (potentially different) multi-level structures in the response variables. We do not assume any particular structure for $\mathbf{G}$, though $\mathbf{b}$ may typically be organized so that $\mathbf{G}$ is block diagonal, as discussed in Section 3.2. Also assume that after conditioning on the random effect vectors, $\mathbf{y}_{(1)}|\mathbf{b}_{(1)}, ..., \mathbf{y}_{(q)}|\mathbf{b}_{(q)}$ are independent. Let $f(\mathbf{y}_{(i)}|\mathbf{b}_{(i)})$ represent the conditional density function for the i-th response variable, determined by the distribution and link function used for that response variable. The joint likelihood may thus be expressed as

$$\begin{aligned} f(\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)}) &= \int f(\mathbf{y}_{(1)}|\mathbf{b}_{(1)}) \ldots f(\mathbf{y}_{(q)}|\mathbf{b}_{(q)}) f(\mathbf{b}) \mathrm{d}\mathbf{b} \\ &= \int f(\mathbf{y}_{(1)}|\mathbf{b}) \ldots f(\mathbf{y}_{(q)}|\mathbf{b}) f(\mathbf{b}) \mathrm{d}\mathbf{b}. \end{aligned} \tag{1}$$

The l-th observation for response variable i, denoted $y_{l(i)}$, has expected value $\mu_{l(i)}$, where for some link function g, $\eta_{l(i)} = g(\mu_{l(i)}) = \mathbf{x}_{l(i)}\boldsymbol{\beta}_{(i)} + \mathbf{z}_{l(i)}\mathbf{b}$. Thus $\mathbf{x}_{l(i)}$ represents the l-th row of the design matrix $\mathbf{X}_{(i)}$ for response variable i, and $\boldsymbol{\beta}_{(i)}$ is the vector of fixed effects for response variable i. The vector $\mathbf{z}_{l(i)}$ has length equal to that of $\mathbf{b}$, and contains a zero corresponding to elements of $\mathbf{b}$ other than $\mathbf{b}_{(i)}$. This definition of $\mathbf{z}_{l(i)}$ streamlines our discussion of the estimation routine. We assume the matrices $\mathbf{X}_{(i)}$ are full rank, but we do not assume any special structure for the matrix $\mathbf{Z}_{(i)}$ which is formed with rows $\mathbf{z}_{l(i)}$. The matrices $\mathbf{Z}_{(i)}$ will typically be sparse, and we capitalize on this sparseness in programming the algorithm to make it applicable to large data sets, as discussed in Section 6.2.

For a normally distributed response $\mathbf{y}_{(i)}$, the error terms are modeled in a covariance matrix $\mathbf{R}_{(i)} = cov(\mathbf{y}_{(i)}|\mathbf{b})$. This matrix may range in structure from a multiple of an identity matrix to an unstructured matrix. However, for single parameter distributions from the exponential family the variance is a function of the mean and not explicitly modeled in a GLMM (Stroup, 2012).

2.2. **Example: Poisson and Binary Responses.** The joint model in Section 2.1 can accommodate a wide variety of situations. In order to make some of the concepts more concrete, we now provide an example using sports teams that we will follow throughout the paper. Suppose that there are p teams that play in n games. Two types of responses are recorded for each of the g games, $g = 1, \ldots, n$: team scores $\mathbf{y}_{(1)}$, and binary win/loss indicators $\mathbf{y}_{(2)}$. We consider sports for which ties are necessarily settled in overtime. Typically, teams are ranked using either the scores or the win-loss records (Karl, 2012); the multiresponse model (1) allows rankings to be done using both. Annis and Craig (2005) consider a similar approach for ranking college football teams.

We model the $2n$ team scores contained in $\mathbf{y}_{(1)}$ using a Poisson distribution. We assign 1 to the variable $y_{g(2)}$ with a "home team" win in game g and 0 otherwise, producing a vector $\mathbf{y}_{(2)} = (y_{1(2)}, \ldots, y_{n(2)})'$. For neutral site games, the "home team" may be chosen arbitrarily: in this example, we only use this distinction to keep track of the win/loss modeled in $y_{g(2)}$.

In a game between teams A and B, the probability of a victory for team A is modeled as a function of the difference between "win propensity ratings" for A and B. Team A's score is modeled as a function of the difference between team A's "offensive rating" and team B's "defensive rating." Team B's score is modeled similarly. The offensive, defensive, and win propensity ratings of the j-th team for $j = 1, \ldots, p$ are assumed to be random effects b_j^o, b_j^d, and b_j^w with $\mathbf{b}_j = (\mathrm{b}_j^o, \mathrm{b}_j^d, \mathrm{b}_j^w)' \sim N_3(\mathbf{0}, \mathbf{G}^*)$, where $\mathbf{G}^*$ is an unstructured covariance matrix and p represents the number of teams being ranked. In addition, $\mathbf{b} \sim N(\mathbf{0}, \mathbf{G})$, where $\mathbf{b} = (\mathbf{b}_1', \ldots, \mathbf{b}_p')'$ and $\mathbf{G}$ is block diagonal with p copies of $\mathbf{G}^*$. To match the notation of Section 2,

$\mathbf{b}_{(1)} = (\mathrm{b}_1^o, \mathrm{b}_1^d, \mathrm{b}_2^o, \mathrm{b}_2^d, \ldots, \mathrm{b}_p^o, \mathrm{b}_p^d)'$, $\mathbf{b}_{(2)} = (\mathrm{b}_1^w, \mathrm{b}_2^w, \ldots, \mathrm{b}_p^w)'$, $n_{(1)} = 2n$, and $n_{(2)} = n$. Notice how the components of $\mathbf{b}_{(2)}$ are folded into those of $\mathbf{b}_{(1)}$ when forming $\mathbf{b}$ in order to ensure that $\mathbf{G}$ is block diagonal.

Assume the conditional team scores $\mathbf{y}_{(1)}|\mathbf{b} = \mathbf{y}_{(1)}|\mathbf{b}_{(1)}$ follow a Poisson distribution, and that the conditional outcomes $\mathbf{y}_{(2)}|\mathbf{b} = \mathbf{y}_{(2)}|\mathbf{b}_{(2)}$ follow a Bernoulli distribution. The fixed effect $\beta_{(1)}$ for the score response models the mean score, and $\mathbf{X}_{(1)}$ is a column of 1's. We consider a parsimonious model with a single mean score. Richer models could explore differences in mean home team and away team scores. As such, no fixed effects are modeled for $\mathbf{y}_{(2)}$.

The random effects design matrix $\mathbf{Z}^*_{(1)}$ is a $2n \times 2p$ matrix that tracks which team's offense was matched against which team's defense to produce each of the $2n$ scores. If $y_{l(1)}$ is the score compiled by the offense of team A against the defense of team B, then $\mathbf{z}^*_{l(1)}$ contains 1 in the position corresponding to the position of the offensive effect of team A, b^o_A, in $\mathbf{b}_{(1)}$, and -1 in the position corresponding to the position of the defensive effect b^d_B of team B, producing a partially crossed design. $\mathbf{Z}^*_{(1)}$ is then expanded into a $2n \times 3p$ matrix $\mathbf{Z}_{(1)}$ by inserting column of zeros into every third column, corresponding to the positions of $\mathbf{b}^w$ in $\mathbf{b}$. Likewise, $\mathbf{Z}^*_{(2)}$ is an $n \times p$ matrix that indicates which teams played each other in each of the n games. Each row $\mathbf{z}^*_{g(2)}$ contains 1 in the column corresponding to the "home team" and -1 in the column corresponding to the "away team". $\mathbf{Z}^*_{(2)}$ is expanded into a $n \times 3p$ matrix $\mathbf{Z}_{(2)}$ by inserting column of zeros in positions corresponding to the locations of of $\mathbf{b}^o$ and $\mathbf{b}^d$ in $\mathbf{b}$.

To illustrate the construction of these matrices, consider a simple example with two teams: A and B. Suppose team A wins a home game against team B by a score of 17 to 10 and then loses a game while visiting team B by a score of 21 to 24. Then $\mathbf{y}_{(1)} = (17, 10, 21, 24)'$ and $\mathbf{y}_{(2)} = (1, 1)'$. Again letting o, d, w represent the offensive, defensive, and win propensity effects, the random effects structure produced by this example is $\mathbf{b} = \left(\mathrm{b}^o_A, \mathrm{b}^d_A, \mathrm{b}^w_A, \mathrm{b}^o_B, \mathrm{b}^d_B, \mathrm{b}^w_B\right)'$,

$$\mathbf{G}^* = \begin{pmatrix} var(o) & cov(o,d) & cov(o,w) \\ cov(d,o) & var(d) & cov(d,w) \\ cov(w,o) & cov(w,d) & var(w) \end{pmatrix},$$

$$\mathbf{G} = \begin{pmatrix} \mathbf{G}^* & \mathbf{0} \\ \mathbf{0} & \mathbf{G}^* \end{pmatrix},$$

$$\mathbf{Z}_{(1)} = \begin{pmatrix} 1 & 0 & 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & 1 & 0 & 0 \end{pmatrix},$$

$$\mathbf{Z}_{(2)} = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & -1 \\ 0 & 0 & -1 & 0 & 0 & 1 \end{pmatrix}.$$

The likelihood functions for the scores (using a log link) and outcomes (using a probit link, with Φ representing the standard normal cumulative distribution function) are

$$f(\mathbf{y}_{(1)}|\mathbf{b}) = \prod_{l=1}^{2n} \left[\frac{1}{y_{l(1)}!} \exp\left\{ y_{l(1)}(\beta_{(1)} + \mathbf{z}'_{l(1)}\mathbf{b}) \right\} \exp\left\{ -\exp\left[\beta_{(1)} + \mathbf{z}'_{l(1)}\mathbf{b} \right] \right\} \right] \tag{2}$$

$$f(\mathbf{y}_{(2)}|\mathbf{b}) = \prod_{g=1}^{n} \left[\Phi\left\{ (-1)^{1-y_{g(2)}}\, \mathbf{z}'_{g(2)}\mathbf{b} \right\} \right]. \tag{3}$$

Typically, teams would be ranked by maximizing only one of the likelihoods, (2) or (3). The joint likelihood function combines both to produce

$$L(\beta_{(1)}, \mathbf{G}) = \int \cdots \int f(\mathbf{y}_{(1)}|\mathbf{b}) f(\mathbf{y}_{(2)}|\mathbf{b}) f(\mathbf{b}) \mathrm{d}\mathbf{b}. \tag{4}$$

where $f(\mathbf{b})$ is the density of $\mathbf{b} \sim N(\mathbf{0}, \mathbf{G})$. The centerpiece of the joint model is the pair of unique off-diagonal covariance terms between $\mathbf{b}_{(1)}$ and $\mathbf{b}_{(2)}$ in the $\mathbf{G}$ matrix, referred to as $cov(w,o)$ and $cov(w,d)$ in the example above. To illustrate: if offensive ratings are highly correlated with the win propensity ratings, then the teams with the best offenses would have larger win propensity ratings in the joint model than they would if only (3) were maximized. If these covariance terms were constrained to 0, the resulting model fit would be equivalent to the one obtained by modeling the two responses independently.

2.3. **Computational Methods for GLMMs.** The likelihood in (1) contains an integral that potentially has the same dimension as the number of random effects. While the marginal likelihood function in completely hierarchical models may be factored into a product of low dimensional integrals and evaluated by quadrature methods (Evans and Swartz, 1995), models with a potentially non-nested structure (e.g. multimembership or cross-classified) produce likelihood functions that must be expressed as a single, high-dimensional, intractable integral. Quadrature methods are infeasible for high-dimensional integrals.

Markov Chain Monte Carlo (MCMC) methods are frequently used in a Bayesian approach for these problems. These, however, often require a subjective prior distribution, and the estimates for random effects models can be very sensitive to the choice of prior distribution, as in Mariano et al. (2010). Judging convergence for these models can be challenging (de Leeuw and Meijer, 2008, p. 357). A method for computing ML estimates avoids the need for a prior distribution, and provides a useful check on sensitivity to the prior even when a Bayesian approach is preferred.

Asymptotic methods such as pseudo-likelihood (PL) linearization (Wolfinger and O'Connell, 1993), penalized quasi-likelihood (PQL, Breslow and Clayton 1993), first-order Laplace approximations, and fully exponential Laplace approximations (Tierney and Kadane, 1986; Tierney et al., 1989) avoid the problems of MCMC methods and the need for prior distributions, though potentially at the cost of accuracy. Pinheiro and Bates (1995) compare PQL, the first-order Laplace approximation, Gaussian quadrature, adaptive Gaussian quadrature (AGQ), and importance sampling, and conclude that the Laplace approximation and AGQ approximations gives the "best mix of efficiency and accuracy." Raudenbush et al. (2000) propose an alternative method for nested GLMMs that uses a sixth-order Laplace approximation. Pinheiro and Chao (2006) provide a scalable Laplacian and AGQ methods for multilevel models with an arbitrary number of nested effects, noting the need for extension to non-nested models.

Most of the work described above uses Newton-Raphson (NR) to find the parameter values maximizing the likelihood. Alternatively, an EM algorithm may be used to estimate the model parameters (Dempster et al., 1977; McLachlan and Krishnan, 2008). EM algorithms treat the random effects as missing data: the E step calculates the conditional expectation of the complete-data likelihood, given the observed data and current parameter estimates, and the M step maximizes the conditional expectation of the complete data likelihood (Laird and Ware, 1982). Although Lindstrom and Bates (1988) highlight advantages of NR methods over the EM algorithm, including quadratic as opposed to linear convergence rates, the EM algorithm provides a viable solution in situations where NR algorithms may fail. For some applications of correlated random effects models, strong correlations in the random effects covariance matrix cause difficulty for Newton methods (Karl et al., 2013b). The EM algorithm provides a stable maximization routine even when presented with nearly perfectly correlated effects.

When estimating nonlinear mixed models with an EM algorithm, the intractable integral in the marginal likelihood produces a similarly intractable integral in the E step. Steele (1996) uses a fully exponential Laplace approximation in the E step of an EM algorithm to estimate GLMMs with uncorrelated random effects, and Rizopoulos et al. (2009) apply this approach to a nested design and relax the assumption of a diagonal random effects covariance matrix. Their estimation routine is available via the R (R Core Team, 2013) package JM (Rizopoulos, 2010).

We use the EM algorithm to compute maximum likelihood estimates of the parameters, using fully exponential and first-order Laplace approximations for the integrals in the E step. The innovations in this approach are (1) a uniform method for the E step that depends on the derivatives of the conditional log likelihood function, which allows the method to be applied to other GLMM settings merely by changing those derivatives, (2) using the sparseness of the random effects design matrices to keep computations manageable, and (3) describing how redundant calculations may be identified and omitted in the presence of non-nested random effects.

3. Estimation Procedure

Let $\mathbf{\Psi}$ be a vector of the unique model parameters. We will refer to $f(\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)}; \mathbf{\Psi})$ as the observed data density function and $f(\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)}, \mathbf{b}; \mathbf{\Psi}) = f(\mathbf{y}_{(1)}|\mathbf{b}; \mathbf{\Psi}) \cdots f(\mathbf{y}_{(q)}|\mathbf{b}; \mathbf{\Psi}) f(\mathbf{b}; \mathbf{\Psi})$ as the complete data density function. Given initial values for the parameters and the random effects, the EM algorithm (Dempster et al., 1977; Laird and Ware, 1982) alternates between an expectation (E) step and a maximization (M) step. At iteration $(k+1)$, the E step calculates the conditional expectation of the complete data log-likelihood, given the observed data, $(\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)})$, and parameter estimates obtained in the k-th step,

$\boldsymbol{\Psi}^{(k)}$. That is, the E step computes

$$Q(\boldsymbol{\Psi};\boldsymbol{\Psi}^{(k)}) = \int \left\{ \left[\sum_{i=1}^{q} \log f\left(\mathbf{y}_{(i)}|\mathbf{b};\boldsymbol{\Psi}\right)\right] + \log f\left(\mathbf{b};\boldsymbol{\Psi}\right) \right\} f(\mathbf{b}|\mathbf{y}_{(1)},\ldots,\mathbf{y}_{(q)};\boldsymbol{\Psi}^{(k)})\mathrm{d}\mathbf{b}. \tag{5}$$

The M step then maximizes $Q(\boldsymbol{\Psi};\boldsymbol{\Psi}^{(k)})$ with respect to $\boldsymbol{\Psi}$, resulting in the updated parameter vector $\boldsymbol{\Psi}^{(k+1)}$ satisfying

$$\int \frac{\partial}{\partial \boldsymbol{\Psi}} \left\{ \left[\sum_{i=1}^{q} \log f\left(\mathbf{y}_{(i)}|\mathbf{b};\boldsymbol{\Psi}\right)\right] + \log f(\mathbf{b};\boldsymbol{\Psi}) \right\} f(\mathbf{b}|\mathbf{y}_{(1)},\ldots,\mathbf{y}_{(q)};\boldsymbol{\Psi}^{(k)})\mathrm{d}\mathbf{b}\Big|_{\boldsymbol{\Psi}=\boldsymbol{\Psi}^{(k+1)}} = \mathbf{0}, \tag{6}$$

Assuming that the order of differentiation and integration may be exchanged – which is true for the exponential family (Lehmann and Romano, 2010) – the expression on the left side of (6) is equivalent to the observed data score vector $S(\boldsymbol{\Psi}) = (\partial/\partial\boldsymbol{\Psi})\, l(\boldsymbol{\Psi};\mathbf{y}_{(1)},\ldots,\mathbf{y}_{(q)})$ (Louis, 1982; McLachlan and Krishnan, 2008). An EM algorithm is used to develop an efficient routine for calculating maximum likelihood estimates of multiple membership linear mixed models (LMMs) under an assumption of missing at random (MAR) by Karl et al. (2012, 2013b). When at least one response is non-normally distributed, the integral in $Q(\boldsymbol{\Psi};\boldsymbol{\Psi}^{(k)})$ has no closed form solution. Because we consider potentially non-nested designs, the dimension of the integral is equal to the length of $\mathbf{b}$. For GLMMs, Steele (1996) proposes approximating the intractable integrals in (5) with a fully exponential Laplace approximation (Tierney et al., 1989). Rizopoulos et al. (2009) use this method to fit a nested shared parameter model for a normal and a time-to-dropout process. To estimate the parameters of the correlated random effects model (1), we generalize the approach for non-nested LMMs proposed by Karl et al. (2013b) by incorporating elements of the E step approximation used by Rizopoulos et al. (2009). The fully exponential Laplace approximation produces computationally expensive corrections to the first-order Laplace approximation. These corrections are not particularly burdensome for the nested design of Rizopoulos et al. (2009) since, in that application, only around 15 thousand such corrections need to be calculated during each iteration. Our aim in this paper is to present the calculation of the model terms for a non-nested design and to discuss computational methods that are required in situations where a naive approach would lead to the calculation of millions of corrections per iteration.

3.1. **Derivatives of the Conditional Log-Likelihood(s).** Let $\eta_{l(i)} = \mathbf{x}'_{l(i)}\boldsymbol{\beta}_{(i)} + \mathbf{z}'_{l(i)}\mathbf{b}$ be the linear predictor (from the GLMM) for the l-th observation for response $\mathbf{y}_{(i)}$. The EM algorithm requires derivatives of $\log f(\mathbf{y}_{(i)}|\mathbf{b})$ with respect to $\mathbf{b}$. We use the fact that $\eta_{l(i)}$ is linear in $\mathbf{b}$ to simplify this calculation. By the chain rule,

$$\frac{\partial \log f(\mathbf{y}_{(i)}|\mathbf{b})}{\partial \mathbf{b}} = \sum_{l=1}^{n_i} \frac{\partial \log f(\mathbf{y}_{(i)}|\mathbf{b})}{\partial \eta_{l(i)}} \cdot \mathbf{z}_{l(i)} \tag{7}$$

A total of four derivatives of $\log f(\mathbf{y}_{(i)}|\mathbf{b})$ with respect to $\eta_{l(i)}$ will be required. Appendix A lists these derivatives for Poisson, binary, and normal distributions.

3.2. **The M step.** The M step of the EM algorithm maximizes the conditional expectation of the complete data log-likelihood. Let $\widetilde{\mathbf{b}} = \mathrm{E}[\mathbf{b}|\mathbf{y}_{(1)},\ldots,\mathbf{y}_{(q)};\boldsymbol{\Psi}]$ and $\widetilde{\mathbf{v}} = \mathrm{var}[\mathbf{b}|\mathbf{y}_{(1)},\ldots,\mathbf{y}_{(q)};\boldsymbol{\Psi}]$ represent the conditional expectation and variance, respectively, of $\mathbf{b}$. These quantities are calculated in the E step and remain fixed during the M step.

3.2.1. *Fixed Effects.* Assuming the fixed effects for each response are distinct, the observed data score function for $\boldsymbol{\beta}_{(i)}$ is

$$\begin{aligned} S(\boldsymbol{\beta}_{(i)}) &= \int \frac{\partial \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \boldsymbol{\beta}_{(i)}} f(\mathbf{b}|\mathbf{y}_{(1)},\ldots,\mathbf{y}_{(q)})\mathrm{d}\mathbf{b} \\ &= \sum_{l=1}^{n_i} \mathbf{x}_{l(i)} \int \frac{\partial \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}} f(\mathbf{b}|\mathbf{y}_{(1)},\ldots,\mathbf{y}_{(q)})\mathrm{d}\mathbf{b} \end{aligned} \tag{8}$$

When $\mathbf{y}_{(i)}$ is normally distributed with error covariance matrix $\mathbf{R}_{(i)}$, the solution for $S(\widehat{\boldsymbol{\beta}}_{(i)}) = 0$ is

$$\widehat{\boldsymbol{\beta}}_{(i)} = \left(\mathbf{X}'_{(i)}\mathbf{R}^{-1}_{(i)}\mathbf{X}_{(i)}\right)^{-1} \mathbf{X}'_{(i)}\mathbf{R}^{-1}_{(i)}\left(\mathbf{y}_{(i)} - \mathbf{Z}_{(i)}\widetilde{\mathbf{b}}\right).$$

By contrast, there is no such closed form solution when $\mathbf{y}_{(i)}$ is non-normally distributed since the integrals in (8) are intractable. The approximations to the integrals in (8) will be obtained in the E step. $\widehat{\boldsymbol{\beta}}_{(i)}$

may then be found by Newton-Raphson: the Hessian $\mathcal{H}(\widehat{\boldsymbol{\beta}}_{(i)})$ is calculated by applying a central difference approximation to $S(\boldsymbol{\beta}_{(i)})$ at $\widehat{\boldsymbol{\beta}}_{(i)}$. Iteration $p+1$ of the approximation yields

$$\widehat{\boldsymbol{\beta}}_{(i)}^{p+1} = \widehat{\boldsymbol{\beta}}_{(i)}^{p} - \mathcal{H}(\widehat{\boldsymbol{\beta}}_{(i)}^{p})^{-1} S(\widehat{\boldsymbol{\beta}}_{(i)}^{p}).$$

With a recommendation of $\alpha = 10^{-8}$, iterations continue until

$$S(\widehat{\boldsymbol{\beta}}_{(i)}^{p})\,'\mathcal{H}(\widehat{\boldsymbol{\beta}}_{(i)}^{p})^{-1} S(\widehat{\boldsymbol{\beta}}_{(i)}^{p}) < \alpha. \tag{9}$$

3.2.2. *The Random Effects Covariance Matrix.* The M step update for $\mathbf{G}$ does not depend on the response distributions included in (1). Throughout Section 3.2.2, assume that U independent groups of random effects (e.g. levels in a multilevel model) are modeled across combinations of the q response variables. Within each group u, $u = 1, \ldots, U$, V_u variance components are modeled with M_u random effects per variance component. These effects are correlated within each group. In the context of the Poisson-binary model in Section 2.2, $U = 1$, $q = 2$, $V_1 = 3$, and $M_1 = p$. An additional game-level effect could be added to the Poisson model for scores to account for intra-game correlation, in which case $U = 2$, $V_2 = 1$, and $M_2 = n$.

Let $\mathbf{b}^{(u,1)}, \ldots, \mathbf{b}^{(u,M_u)}$ be the length-$(V_u \times M_u)$ sub-vectors of $\mathbf{b}$ representing the effects in group $u = 1, \ldots, U$. Then $\mathbf{b}$ may be arranged so that

$$\begin{aligned}\mathbf{b} &= \left(\mathbf{b}^{(1,1)}, \ldots, \mathbf{b}^{(1,M_1)}, \mathbf{b}^{(2,1)}, \ldots, \mathbf{b}^{(2,M_2)}, \ldots, \mathbf{b}^{(U,1)}, \ldots, \mathbf{b}^{(U,M_U)}\right) \\ \mathbf{G} &= \text{block diagonal}\left(\boldsymbol{\Gamma}_1, \ldots, \boldsymbol{\Gamma}_1, \ldots, \boldsymbol{\Gamma}_2, \ldots, \boldsymbol{\Gamma}_2, \ldots, \boldsymbol{\Gamma}_U, \ldots, \boldsymbol{\Gamma}_U\right).\end{aligned}$$

$\mathbf{G}$ contains M_u copies each of the $V_u \times V_u$ matrix $\boldsymbol{\Gamma}_u = cov(\mathbf{b}^{(u,j)})$, for $j = 1, \ldots, M_u$. Let the components of $\widetilde{\mathbf{b}}$ corresponding to $\mathrm{E}[\mathbf{b}^{(u,j)}|\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)}; \boldsymbol{\Psi}]$ be denoted $\widetilde{\mathbf{b}}^{(u,j)}$ for $u = 1, \ldots, U$ and $j = 1, \ldots, M_u$. Likewise, let $\widetilde{\mathbf{v}}^{(u,j)}$ represent the block of the matrix $\widetilde{\mathbf{v}}$ corresponding to $\mathrm{E}[\mathbf{b}^{(u,j)}(\mathbf{b}^{(u,j)})'|\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)}; \boldsymbol{\Psi}]$. When the $\boldsymbol{\Gamma}_u$ are unstructured, a generalization of steps illustrated by Karl et al. (2013b) shows

$$\widehat{\boldsymbol{\Gamma}}_u = \frac{1}{M_u} \sum_{j=1}^{M_u} \left\{ \widetilde{\mathbf{v}}^{(u,j)} + \widetilde{\mathbf{b}}^{(u,j)} \left(\widetilde{\mathbf{b}}^{(u,j)}\right)' \right\} \tag{10}$$

Alternative structures may be imposed on the $\boldsymbol{\Gamma}_u$ using the same approach as that demonstrated for an AR(1) (first-order autoregressive) $\mathbf{R}$ matrix in the next section. This is a general approach: all that is required is the first derivative of $\mathbf{G}$ with respect to each unique parameter in the matrix. This also allows for the possibility of a non-block-diagonal $\mathbf{G}$. For the E step, the groupings of random effects (e.g. levels of a multilevel model) do not need to be tracked explicitly, since this information is contained in the structure of $\mathbf{G}$ and components of the $\mathbf{Z}_{(i)}$ matrices.

3.2.3. *Error Covariance Matrix for a Normal Response.* For a normally distributed response $\mathbf{y}_{(i)}$, the error covariance matrix $\mathbf{R}_{(i)}$ may take on a number of possible structures. Karl et al. (2013b) provide the calculations for an unstructured covariance matrix in the presence of incomplete data. We demonstrate the calculations of an AR(1) matrix assuming complete data, $\mathbf{y}$. For this single response, we will drop the subscript (i) from the model terms defined previously. Suppose that $\mathbf{y} \sim N_n(\mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{b}, \mathbf{R})$ and the (k, d) element of $\mathbf{R}$ is $\sigma^2\rho^{|k-d|}$. We will derive the score equations for σ^2 and ρ. The solution $S(\sigma^2, \rho) = \mathbf{0}$ may then be obtained by Newton-Raphson.

$$\begin{aligned} S(\rho) &= \int \frac{\partial}{\partial \rho} \left[\log\left(|\mathbf{R}|^{-1/2}\right) - \frac{1}{2}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\mathbf{b})' \mathbf{R}^{-1} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\mathbf{b}) \right] f(\mathbf{b}|\mathbf{y}; \boldsymbol{\Psi}) \mathrm{d}\mathbf{b} \\ &= -\frac{1}{2} \int \left[\mathrm{tr}\left(\mathbf{R}^{-1} \frac{\partial \mathbf{R}}{\partial \rho}\right) - (\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\mathbf{b})' \mathbf{R}^{-1} \frac{\partial \mathbf{R}}{\partial \rho} \mathbf{R}^{-1} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\mathbf{b}) \right] f(\mathbf{b}|\mathbf{y}; \boldsymbol{\Psi}) \mathrm{d}\mathbf{b} \\ &= -\frac{1}{2} \left[\mathrm{tr}\left(\mathbf{R}^{-1} \frac{\partial \mathbf{R}}{\partial \rho}\right) - \left(\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\widetilde{\mathbf{b}}\right)' \mathbf{R}^{-1} \frac{\partial \mathbf{R}}{\partial \rho} \mathbf{R}^{-1} \left(\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\widetilde{\mathbf{b}}\right) - \mathrm{tr}\left(\mathbf{Z}'\mathbf{R}^{-1} \frac{\partial \mathbf{R}}{\partial \rho} \mathbf{R}^{-1} \mathbf{Z}\widetilde{\mathbf{v}}\right) \right] \end{aligned} \tag{11}$$

For some patterned matrices or for data sets with incomplete data, $\mathbf{R}^{-1}$ must be calculated numerically. Furthermore, when some elements of $\mathbf{y}$ are missing, the score function is the sum of values calculated from (11) over each dropout pattern in the data set (Karl et al., 2013b). However, for a complete AR(1) matrix $\mathbf{R}$, $\mathbf{R}^{-1}$ has a simple tridiagonal expression (Demidenko, 2013) and thus

$$\mathrm{tr}\left(\mathbf{R}^{-1} \frac{\partial \mathbf{R}}{\partial \rho}\right) = \frac{-2\rho(n-1)}{1-\rho^2}.$$

Likewise,

$$S\left(\sigma^2\right) = -\frac{1}{2}\left[\frac{n}{\sigma^2} - \left(\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\widetilde{\mathbf{b}}\right)' \mathbf{R}^{-1}\frac{\partial \mathbf{R}}{\partial \sigma^2}\mathbf{R}^{-1}\left(\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\widetilde{\mathbf{b}}\right) - \operatorname{tr}\left(\mathbf{Z}'\mathbf{R}^{-1}\frac{\partial \mathbf{R}}{\partial \sigma^2}\mathbf{R}^{-1}\mathbf{Z}\widetilde{\mathbf{v}}\right)\right]$$

Finally, if $\mathbf{R} = \sigma^2 I_n$, then

$$\widehat{\sigma}^2 = \frac{1}{n}\left[\left(\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\widetilde{\mathbf{b}}\right)' \left(\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}\widetilde{\mathbf{b}}\right) + \operatorname{tr}\left(\mathbf{Z}'\mathbf{Z}\widetilde{\mathbf{v}}\right)\right]. \tag{12}$$

3.3. **The E step.** Calculation of the components of the observed data score vector requires the conditional mean, $\widetilde{\mathbf{b}}$, and the conditional variance, $\widetilde{\mathbf{v}}$, of $f\left(\mathbf{b}|\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)}; \boldsymbol{\Psi}\right)$, as well as the conditional expectations appearing in (8). Letting $\{E\left[H(\mathbf{b})\right]\}_k$ denote the (k)-th component of the vector $E\left[H(\mathbf{b})\right]$, the M step updates require

$$\left\{\mathrm{E}\left[H(\mathbf{b})|\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)}; \boldsymbol{\Psi}\right]\right\}_k = \int \{H(\mathbf{b})\}_k f(\mathbf{b}|\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)}; \boldsymbol{\Psi})\mathrm{d}\mathbf{b}$$

for all k, where $H(\cdot)$ is a function of the random effects, and $\boldsymbol{\Psi}$ is fixed at its value from the previous iteration. The M step updates for $\boldsymbol{\beta}_{(i)}, \mathbf{G}$ and $\mathbf{R}$ require $H(\mathbf{b}) = \mathbf{b}$ and $\widetilde{\mathbf{v}} = \operatorname{var}[\mathbf{b}|\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)}; \boldsymbol{\Psi}]$. When response variable i is non-normally distributed, the M step update for $\boldsymbol{\beta}_{(i)}$ requires

$$H(\mathbf{b}) = \frac{\partial \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}}$$

for all $l \in \{1, \ldots, n_{(i)}\}$. To solve these high-dimensional integration problems, we follow the examples of Steele (1996) and Rizopoulos et al. (2009) and use the fully exponential Laplace approximation of Tierney et al. (1989), approximating the cumulant-generating function $\log\left[E\left(\exp\left[\boldsymbol{c}'H(\mathbf{b})\right]\right)\right]$ at the mode $\widehat{\mathbf{b}} = \widehat{\mathbf{b}}^{(\mathbf{0})}$, where

$$\widehat{\mathbf{b}}^{(\boldsymbol{c})} = \operatorname{argmax}_{\mathbf{b}}\left\{\log\left[f\left(\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)}, \mathbf{b}\right) + \boldsymbol{c}'H\left(\mathbf{b}\right)\right]\right\}.$$

The mode $\widehat{\mathbf{b}}$ is obtained by Newton-Raphson, using the same convergence criterion as (9):

$$\widehat{\mathbf{b}}^{\mathrm{p}+1} = \widehat{\mathbf{b}}^{\mathrm{p}} - \left(\boldsymbol{\Sigma}^{\mathrm{p}}\right)^{-1}\boldsymbol{L}\left(\widehat{\mathbf{b}}^{\mathrm{p}}\right)$$

where superscript "p" is the iteration counter. Via matrix calculus (Magnus and Neudecker, 1999; Harville, 2008),

$$\begin{aligned}\boldsymbol{L}(\mathbf{b}) &= -\frac{\partial}{\partial \mathbf{b}}\left\{\left[\sum_{i=1}^{q}\log f\left(\mathbf{y}_{(i)}|\mathbf{b}; \boldsymbol{\Psi}\right)\right] + \log\left\{f\left(\mathbf{b}\right)\right\} + \boldsymbol{c}'H(\mathbf{b})\right\}|_{\boldsymbol{c}=\mathbf{0}} \\ &= -\sum_{i=1}^{q}\sum_{l=1}^{n_i}\frac{\partial \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}}\mathbf{z}_{l(i)} + \mathbf{G}^{-1}\mathbf{b} \qquad (13)\end{aligned}$$

and $\boldsymbol{\Sigma}^{w} = \boldsymbol{\Sigma}^{(\boldsymbol{c})}|_{(\boldsymbol{c},\mathbf{b})=(\mathbf{0},\widehat{\mathbf{b}}^{\mathrm{w}})}$, with

$$\begin{aligned}\boldsymbol{\Sigma}^{(c)} &= -\frac{\partial^2}{\partial \mathbf{b}\partial \mathbf{b}'}\left\{\left[\sum_{i=1}^{q}\log f\left(\mathbf{y}_{(i)}|\mathbf{b}; \boldsymbol{\Psi}\right)\right] + \log\left\{f\left(\mathbf{b}\right)\right\} + \boldsymbol{c}'H(\mathbf{b})\right\} \\ &= -\sum_{i=1}^{q}\sum_{l=1}^{n_i}\frac{\partial^2 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^2}\mathbf{z}_{l(i)}\mathbf{z}_{l(i)}' + \mathbf{G}^{-1} - \frac{\partial^2 \boldsymbol{c}'H(\mathbf{b})}{\partial \mathbf{b}\partial \mathbf{b}'}. \qquad (14)\end{aligned}$$

For large problems, the algorithm may be sped up by performing only a single NR step in early iterations and then running the algorithm to convergence in later iterations. Although the theory of the EM algorithm requires finding the conditional mode in the E step, this process requires more NR iterations in the first few E steps as the parameter estimates begin to stabilize. The resulting parameter estimates, $\boldsymbol{\Psi}^{(k)}$, after k such abbreviated EM iterations may be viewed as initial conditions for the full EM algorithm. Of course, there is the usual caveat that the EM algorithm only guarantees convergence to a local maximum, and that the sensitivity to the initial values should be explored in each application.

Once the NR algorithm converges to the mode $\widehat{\mathbf{b}}$, the next step is to apply a fully exponential Laplace approximation to $E[\exp\{\boldsymbol{c}'H(\mathbf{b})\}]$. We apply the result of Theorem 2 of Tierney et al. (1989). Using properties of the cumulant-generating function,

$$\begin{aligned}\left\{E\left[H(\mathbf{b})|\mathbf{y}_{(1)},\ldots,\mathbf{y}_{(q)};\boldsymbol{\Psi}\right]\right\}_k &\approx \frac{\partial}{\partial c_k}\left\{\boldsymbol{c}'H\left(\widehat{\mathbf{b}}^{(\mathbf{c})}\right)+\log\left[\det\left(\boldsymbol{\Sigma}^{(\mathbf{c})}\right)^{-1/2}\right]\right\}\Bigg|_{\boldsymbol{c}=\mathbf{0}} \\ &= \boldsymbol{e}_k'H(\widehat{\mathbf{b}})-\frac{1}{2}\mathrm{tr}\left(\boldsymbol{\Sigma}^{-1}\left\{\frac{\partial\boldsymbol{\Sigma}^{(c)}}{\partial c_k}\Bigg|_{(\mathbf{c},\mathbf{b})=(\mathbf{0},\widehat{\mathbf{b}})}\right\}\right), \qquad (15)\end{aligned}$$

where $\boldsymbol{e}_k$ is the vector of zeros with a 1 in the k-th component. The (k,d) element of $\mathrm{var}(\mathbf{b})$, evaluated at $(\boldsymbol{c},\mathbf{b})=(\mathbf{0},\widehat{\mathbf{b}})$, is

$$\left\{\mathrm{var}\left(\mathbf{b}|\mathbf{y}_{(1)},\ldots,\mathbf{y}_{(q)};\boldsymbol{\Psi}\right)\right\}_{kd} \approx \boldsymbol{e}_k'\boldsymbol{\Sigma}^{-1}\boldsymbol{e}_d-\frac{1}{2}\mathrm{tr}\left(\boldsymbol{\Sigma}^{-1}\frac{\partial^2\boldsymbol{\Sigma}}{\partial c_k\partial c_d}-\boldsymbol{\Sigma}^{-1}\frac{\partial\boldsymbol{\Sigma}}{\partial c_d}\boldsymbol{\Sigma}^{-1}\frac{\partial\boldsymbol{\Sigma}}{\partial c_k}\right). \qquad (16)$$

The first-order Laplace approximation consists of only the first terms of (15) and (16) (Kass and Steffey, 1989). The terms involving the trace in both equations are the fully exponential corrections to the first-order Laplace approximation. As discussed in Section 4, for problems where the fully exponential corrections are too expensive to calculate, the first-order approximation may be used instead. Notice that the first-order Laplace approximation only uses the first two derivatives of the conditional log-likelihood. We will see that the fully exponential corrections in (15) and (16) require the third and fourth derivatives, respectively. This suggests that the EBLUPs ($\widetilde{\mathbf{b}}$) from GLMMs using different link functions for the same distribution will show greater differences when the fully exponential corrections are applied.

Calculations of the fully exponential corrections are furnished in Appendix B. These are generalizations of the calculations presented by Rizopoulos et al. (2009). The non-nested structure of our model affects the terms derived inside of the trace functions.

3.4. **Convergence and EM Standard Errors.** The EM algorithm converges to a stationary value of the approximated observed data likelihood as long as the E and M step updates are continuous in the model parameters, $\boldsymbol{\Psi}$, and the parameter space is compact (Wu, 1983). Although the parameter space for $\boldsymbol{\Psi}$ is not compact for our model, this regularity condition is satisfied by a truncation of the parameter space (McCulloch, 1994; Demidenko, 2013). The existence of the derivatives that yield the score functions in the M step guarantees that the M step is continuous with respect to $\boldsymbol{\Psi}$. Finally, the E step functions (15) and (16) are continuous functions of the elements of $\boldsymbol{\Psi}$. Due to the approximation used in the E step, the limit point lies on a surface approximating the likelihood function: the quality of this approximation is discussed in Section 3.5.

EM algorithms have been documented to converge slowly when there is a large number of random effects, or when a variance component is near zero. Lindstrom and Bates (1988) note how the EM algorithm is often outperformed by a Newton-Raphson routine when fitting mixed models. However, NR algorithms tend to step outside of the parameter space when the $\mathbf{G}$ matrix is nearly singular. Karl et al. (2013b) show how the NR routine of SAS GLIMMIX fails in the presence of nearly perfectly correlated random effects, while an EM algorithm converges successfully since the estimated $\mathbf{G}$ matrix is guaranteed to be positive definite after each iteration. It is feasible that applications of the joint model presented in Section 2 will have highly correlated random effects across response variables, especially in situations when assumptions made by shared parameter models (Wu and Carroll, 1988) are valid. Rizopoulos et al. (2009) demonstrate how an EM approach may be converted to a quasi-Newton approach, which could be useful in the absence of strong correlations.

The observed data score vector $S(\boldsymbol{\Psi})$ used in the M step is equal to the the conditional expectation of the complete data score vector (Louis, 1982). As proposed by Jamshidian and Jennrich (2000), it is possible to calculate the observed information matrix at the MLE $\widehat{\boldsymbol{\Psi}}$ using a central difference approximation to the Hessian

$$-\partial S(\boldsymbol{\Psi})/\partial\boldsymbol{\Psi}|_{\boldsymbol{\Psi}=\widehat{\boldsymbol{\Psi}}}\,.$$

3.5. **Approximation Error.** In this section we discuss the degree of error in the fully exponential Laplace approximation presented in Section 3. Theorem 1 of Tierney et al. (1989) demonstrates that, for the

approximations appearing in (15) and (16),

$$E[H] = \hat{E}[H] + O(r^{-2})$$
$$V[H] = \hat{V}[H] + O(r^{-3})$$

where the hat denotes the fully exponential Laplace approximation, and $r = r(\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)})$ is a measure of the size of the data set $(\mathbf{y}_{(1)}, \ldots, \mathbf{y}_{(q)})$ such that $r \to \infty$ as the size of the data set grows (Evans and Swartz, 1995). Applying the Laplace approximation to a non-normal model $f(\mathbf{y}_{(i)}|\mathbf{b})$ will result in approximation error because the $\mathbf{b}_i$ enter the integrand via a nonlinear link function. The approximation error depends on the amount of information in $\mathbf{y}_{(i)}$ and the random effects structure of $f(\mathbf{y}_{(i)}|\mathbf{b})$. If the random effects structure is nested, then r equals the minimum number of observations on any of the components of $\mathbf{b}_i$, along the lines of the work by Vonesh (1996). A numerical study of the performance of the first-order Laplace approximation in nested binary and Poisson models is provided by Joe (2008).

Models for normally distributed responses do not contribute to the approximation error. For illustration, consider the joint modeling of a normal response $\mathbf{y}_{(1)}$ with a non-normal response $\mathbf{y}_{(2)}$. Applying the Laplace approximation to the marginal likelihood,

$$\begin{aligned}
&\iint f(\mathbf{y}_{(1)}|\mathbf{b}_{(1)}) f(\mathbf{y}_{(2)}|\mathbf{b}_{(2)}) f(\mathbf{b}_{(1)}, \mathbf{b}_{(2)}) \mathrm{d}\mathbf{b}_{(1)} \mathrm{d}\mathbf{b}_{(2)} \\
&= \int f(\mathbf{y}_{(2)}|\mathbf{b}_{(2)}) \left\{ \int f(\mathbf{y}_{(1)}|\mathbf{b}_{(1)}) f(\mathbf{b}_{(1)}, \mathbf{b}_{(2)}) \mathrm{d}\mathbf{b}_{(1)} \right\} \mathrm{d}\mathbf{b}_{(2)} \\
&= \int f(\mathbf{y}_{(2)}|\mathbf{b}_{(2)}) I(\mathbf{b}_{(2)}; \mathbf{y}_{(1)}) \mathrm{d}\mathbf{b}_{(2)}. \qquad (17)
\end{aligned}$$

The Laplace approximation is exact for $\int f(\mathbf{y}_{(1)}|\mathbf{b}_{(1)}) f(\mathbf{b}_{(1)}, \mathbf{b}_{(2)}) \mathrm{d}\mathbf{b}_{(1)}$, since $\mathbf{b}_{(1)}$ enters the model linearly (Pinheiro and Bates, 1995). The result, $I(\mathbf{b}_{(2)}; \mathbf{y}_{(1)})$, is the density of a normal distribution with mean and variance depending on the covariance between $\mathbf{b}_{(1)}$ and $\mathbf{b}_{(2)}$. For example, if these effects are uncorrelated – and hence independent, due to their joint normality – then $I(\mathbf{b}_{(2)}; \mathbf{y}_{(1)}) = f(\mathbf{b}_{(2)}) f(\mathbf{y}_{(1)})$.

If the number of random effects in the non-normal model $f(\mathbf{y}_{(i)}|\mathbf{b}_{(i)})$ increases with the sample size, then the dimension of the approximated integrals in the E step also increases with the sample size. This property is not typical for applications of the Laplace approximation; Shun and McCullagh (1995) describe how a modification to the first-order Laplace approximation is needed to retain its usual order of convergence, although they did not study the behavior of the fully exponential approximation, and note that it may not suffer the same extent of deterioration in this setting as the first-order approximation. The results of Shun (1997) show, in an application to the salamander mating data (McCullagh and Nelder, 1989), that the uncorrected first-order Laplace approximation outperforms PQL. This is reasonable since PQL makes an additional approximation to one of the terms in the first-order Laplace approximation. At worst, the first-order and fully exponential Laplace approximations present an improvement over PQL (and thus PL when the scale parameter is fixed) in these settings, even if consistency is not guaranteed.

4. Computation

In order to compute the fully exponential approximation in a reasonable amount of time, we use a few key facts. First, the operations inside of the trace function in (16) contain some elements that are common to all of the components of $\widetilde{\mathbf{v}}$ and do not need to be re-calculated for each component. Secondly, some of the needed computations involve multiplication by sparse matrices, greatly reducing the number of needed arithmetic operations. Finally, not all of the components of $\widetilde{\mathbf{v}}$ will be used in the M step: the unused components may be ignored.

To see why only some of the components of $\widetilde{\mathbf{v}}$ are needed, note that the M step update for $\boldsymbol{\beta}_{(i)}$ does not use $\widetilde{\mathbf{v}}$ and observe how the covariance matrix $\widetilde{\mathbf{v}}$ is used in the M step update of $\boldsymbol{\Gamma}$ in (10). The M step update of $\boldsymbol{\Gamma}$ requires only relatively small block-diagonal portions of $\widetilde{\mathbf{v}}$. In the context of the Poisson-binary sports ranking model, these are 3×3 blocks, corresponding to the non-zero components of $\mathbf{G}$. As a result, only 756 of the 8001 components in the upper triangle of $\widetilde{\mathbf{v}}$ will be used in the application of Section 5.1.

For a normally distributed response i with a diagonal error covariance matrix, the M step update of σ^2 requires

$$\mathrm{tr}\left[\mathbf{Z}'_{(i)} \mathbf{Z}_{(i)} \widetilde{\mathbf{v}}\right] = \sum_{\text{components}} \left[\left(\mathbf{Z}'_{(i)} \mathbf{Z}_{(i)}\right) \circ \widetilde{\mathbf{v}}\right] \qquad (18)$$

FIGURE 1. The shaded pixels (including the blocks on the diagonal) indicate which components of $\widetilde{\mathbf{v}}$ need fully exponential corrections from (16) in the application of Section 5.2.

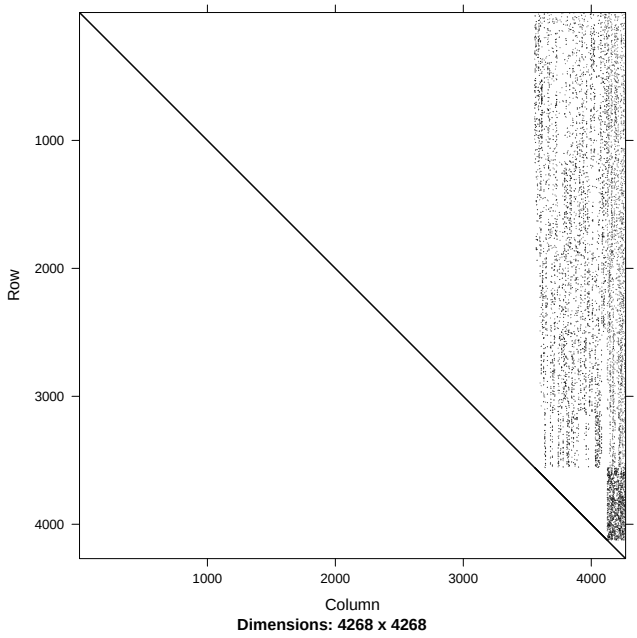

where $\circ$ represents the Hadamard product. Thus, the only components of $\widetilde{\mathbf{v}}$ that are needed in addition to the aforementioned block-diagonal elements are those that have the same indices as non-zero components of the sparse matrix $\mathbf{Z}'_{(i)}\mathbf{Z}_{(i)}$. The plot in Figure 1 indicates which components of $\widetilde{\mathbf{v}}$ require corrections in the application of Section 5.2. Only $14138/9110046 = 0.0016$ of the elements in the upper triangle of $\widetilde{\mathbf{v}}$ need to be calculated. Unfortunately, this feature does not extend to normally distributed responses with a non-diagonal error covariance matrix, $\mathbf{R}$. As is apparent from the score function for ρ in (11), the updates for the components of $\mathbf{R}$ depend on

$$\sum_{\text{components}} \left[\left(\mathbf{Z}'\mathbf{R}^{-1} \frac{\partial \mathbf{R}}{\partial \rho} \mathbf{R}^{-1}\mathbf{Z} \right) \circ \widetilde{\mathbf{v}} \right]$$

Although such models are capable of being estimated by the proposed method with a fully exponential Laplace approximation, the computations are not scalable. For these or other situations where the fully exponential corrections are too expensive to calculate, the first-order Laplace approximation produced by the EM algorithm will still be scalable. Or, as suggested in Section 6, a partial fully exponential correction using only (15) may still be calculated efficiently. The main computational limitation on the first-order Laplace approximation is the memory required to hold $\mathbf{\Sigma}^{-1}$, which is dense. With m representing the length of $\mathbf{b}$, this matrix requires approximately $8m^2$ bytes to store in R (half of that if it is explicitly stored as a symmetric matrix). One gigabyte is necessary for m=11,180, two gigabytes for m=15,811, three gigabytes for m=19,364, etc.

The computational burden of calculating the fully exponential corrections in (16) is tremendous. The dimensions of each of the matrices $\mathbf{\Sigma}^{-1}$, $\partial\mathbf{\Sigma}/\partial c_k$, and $\partial^2\mathbf{\Sigma}/(\partial c_k \partial c_d)$ is $m \times m$. Even ignoring the calculations required to obtain the derivatives $\partial\mathbf{\Sigma}/\partial c_k$ and $\partial^2\mathbf{\Sigma}/(\partial c_k \partial c_d)$, the term inside the trace function of (16) requires approximately $4(2m^3 - m^2)$ calculations. This results in a requirement of just over 6.2×10^{11} floating point operations to calculate the fully exponential correction for a single component of $\widetilde{\mathbf{v}}$ for the example in Section 5.2, if implemented naively. Since, in this example, there are 9.1 million components in the upper-triangle of the symmetric $\widetilde{\mathbf{v}}$ matrix, this implementation would require around 5.7×10^{18} operations per iteration, excluding the operations needed to calculate the derivatives. The computer used for this application runs at 7.2 billion floating point operations per second, and would require around 25 years to execute one iteration of a naively implemented fully exponential correction to the full $\widetilde{\mathbf{v}}$ matrix. We reduce this time to 27 hours by calculating only the necessary corrections and by using properties of the trace such as the one shown in (18).

For smaller problems, it would be possible to calculate the derivatives $\partial\boldsymbol{\Sigma}/\partial c_d$, $d = 1, \ldots, m$ and save them in memory, greatly speeding up the fully exponential approximation. However, this would require too much memory for most applications. As a compromise, we first calculate and store $\partial\boldsymbol{\Sigma}/\partial c_1$. We then calculate all of the corrections that require this matrix, and then repeat for $c_2, \ldots, c_m$. The loop calculating the corrections for these components is an excellent candidate for parallelization, which could be done using one of the packages for parallel computing described in Eugster et al. (2011).

The simulations in Section 6 suggest that using the corrections to $\widetilde{\mathbf{b}}$ inside the trace function of (15), but not those to $\widetilde{\mathbf{v}}$ in the trace function of (16), produces results that are relatively close to those obtained by the fully exponential approximations while requiring only a fraction of the time. This seems reasonable since the M step updates for the $\boldsymbol{\beta}_{(i)}$ are functions of $\widetilde{\mathbf{b}}$, but not $\widetilde{\mathbf{v}}$. Furthermore, the M step update for $\mathbf{G}$ is composed of diagonal blocks from the matrix $\widetilde{\mathbf{v}} + \widetilde{\mathbf{b}}\widetilde{\mathbf{b}}'$. Hence $\widehat{\mathbf{G}}$ does experience at least part of the benefit of the fully exponential corrections when only the corrections to $\widetilde{\mathbf{b}}$ are included. We take advantage of this by first running the first-order Laplace approximation to convergence, then introducing the corrections in the trace function of (15) and again run to convergence. We then treat the resulting parameter estimates as initial values for the fully exponential approximation (using the corrections from the trace functions of both (15) and (16)). This minimizes the number of iterations that require corrections to $\widetilde{\mathbf{v}}$, which is important given the tremendous amount of time taken to calculate these corrections, as seen in Table 3.

We rely heavily on the sparse matrix methods of the Matrix package (Bates and Maechler, 2012). To illustrate, only 13534 components of the 5697×4268 $\mathbf{Z}_{(1)}$ matrix and 3557 components of the 3557×4268 $\mathbf{Z}_{(2)}$ matrix for the application in Section 5.2 are nonzero. This has important implications for the program: for example, the calculation of $\mathbf{Z}'_{(1)}\mathbf{Z}_{(1)}$ in R registers a time of 0 seconds when $\mathbf{Z}_{(1)}$ is stored as a sparse matrix (requiring 0.18 MB), and 57 seconds when it is treated as dense (requiring 194.52 MB). Karl et al. (2013b) describe how SAS runs out of memory when attempting to fit a multiple membership linear mixed model with around 3500 random effects on a machine with 4 GB of RAM.

5. Applications

Two applications are presented in this section to illustrate the computational methods of this paper.

5.1. Sports Ranking Model.

The first application fits the Poisson-binary model in Section 2.2 to the regular season data from the 2012 NCAA FBS Division of college football. This data set contains results from 844 games including 126 teams (lower-division teams are consolidated into a single effect, labeled "fcs"). Table 2 summarizes the median time required per iteration for the first-order Laplace approximation, fully exponential corrections from (15) only, and the complete fully exponential Laplace approximation.

The EBLUPs for the offensive and defensive ratings are plotted in Figure 2. The size of each plotted point is proportional to the EBLUP for the team's win propensity. Table 1 presents the estimated $\mathbf{G}$ covariance matrices from the first-order Laplace, first-order Laplace with corrections from (15), and fully exponential Laplace approximations. The offensive and defensive ratings are positively correlated, suggesting teams with good (bad) offenses tend to also have good (bad) defenses. Furthermore, it appears that these ratings are equally correlated with the win propensity. Notice how the estimates for the variance components increase with the quality of the approximation, and how the estimated components using the correction from (15) are relatively close to those from the fully exponential approximation.

Fitting the Poisson and binary models separately leads to the estimates in the second row of Table 1. The most noticeable difference is that the variance of the win propensity rating increases in the joint model. To a lesser extent, the variance of the offensive and defensive ratings increase as well.

5.2. Model for Continuous Outcome with Missing Data Indicator.

Longitudinal studies commonly finish with missing observations: ignoring them requires an assumption that the observations are missing at random (MAR). When the missingness process at time t depends on the values that would be observed at time t, or on latent variables contributing to the value of the observation at time t, the missing observations are missing not at random (MNAR). In this situation, obtaining the appropriate parameter estimates and standard errors requires consideration of a model for the missingness mechanism (Little and Rubin, 2002).

Karl et al. (2013a) develop a correlated random effects model to explore the sensitivity of the EBLUPs from the "generalized persistence" (GP) value-added model (VAM) (Mariano et al., 2010) to assumptions about the missing data. They present an example using calculus 2 and 3 grades from a large public university, with 5697 grades recorded for 3557 students from 189 and 144 calculus 2 and 3 courses, respectively. In the notation of Section 3.2.2, $U = 3$, $q = 2$, $V_1 = 3$, $M_1 = 189$, $V_2 = 1$, $M_2 = 144$, $V_3 = 1$, $M_3 = 3557$. In

TABLE 1. Estimated **G** covariance matrices from the joint Poisson-binary model for scores and win/loss indicators (first row) and from independent models for the scores and indicators (second row).

First-order	FE: (15) only	Fully Exponential (FE)
$\begin{pmatrix} 0.0917 & 0.0472 & 0.3151 \\ 0.0472 & 0.1022 & 0.3355 \\ 0.3151 & 0.3355 & 1.4732 \end{pmatrix}$	$\begin{pmatrix} 0.0920 & 0.0473 & 0.3215 \\ 0.0473 & 0.1025 & 0.3425 \\ 0.3215 & 0.3425 & 1.5297 \end{pmatrix}$	$\begin{pmatrix} 0.0920 & 0.0473 & 0.3216 \\ 0.0473 & 0.1025 & 0.3425 \\ 0.3216 & 0.3425 & 1.5306 \end{pmatrix}$
$\begin{pmatrix} 0.0908 & 0.0460 & 0 \\ 0.0460 & 0.1015 & 0 \\ 0 & 0 & 0.7614 \end{pmatrix}$	$\begin{pmatrix} 0.0908 & 0.0460 & 0 \\ 0.0460 & 0.1015 & 0 \\ 0 & 0 & 1.1105 \end{pmatrix}$	$\begin{pmatrix} 0.0908 & 0.0460 & 0 \\ 0.0460 & 0.1015 & 0 \\ 0 & 0 & 1.1407 \end{pmatrix}$

TABLE 2. Median time per iteration (seconds) for the 2012 college football data.

Approximation	Binary	Poisson	Poisson-binary
First-order	0.11	0.80	1.20
FE: (15) only	1.07	4.32	8.72
Fully Exponential (FE)	1.79	11.33	33.90

FIGURE 2. College football defensive vs. offensive ratings at the end of the regular season. The size of the plotted point increases with the win propensity rating of each team.

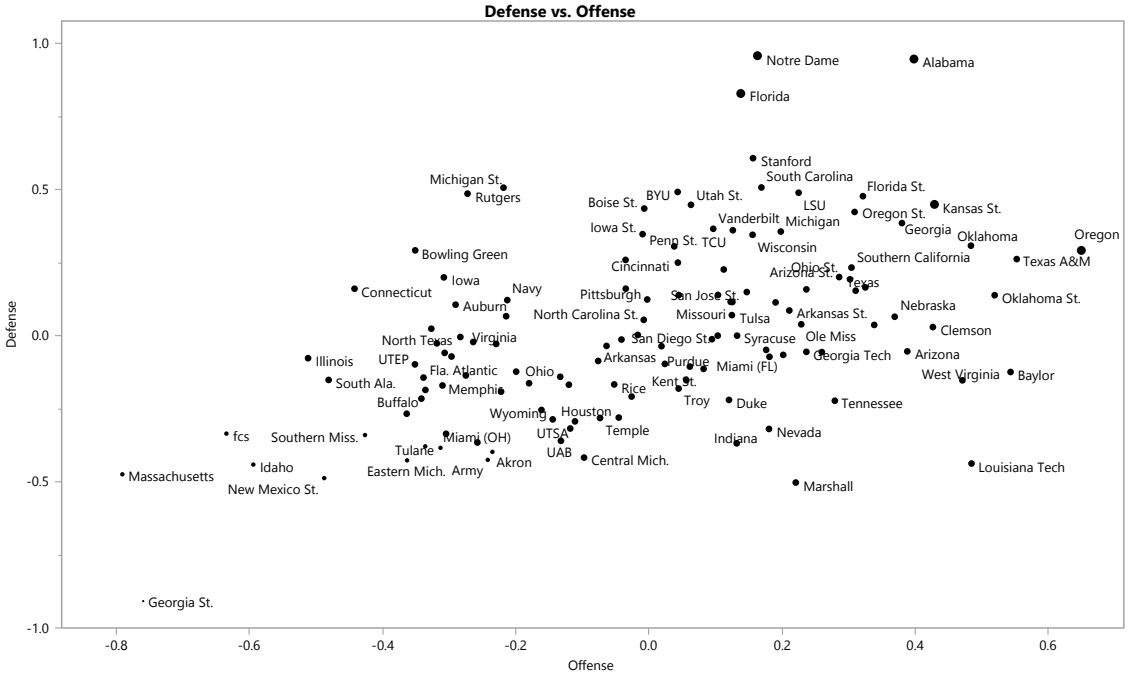

the complete-data GP VAM (Mariano et al., 2010), each calculus 2 teacher has an effect on their students' calculus 2 grades and an effect on their students' calculus 3 grades. Calculus 3 teachers have an effect associated with their students' calculus 3 grades. The two effects of the calculus 2 teachers are allowed to be correlated via the 2×2 covariance matrix $\mathbf{\Gamma}_{teach,1}$. These are assumed to be independent from the effects of the calculus 3 teachers, whose effects are modeled by the variance component $\Gamma_{teach,2}$. Furthermore, a random intercept for each student, with associated variance component Γ_{stu}, allows the student grades for calculus 2 and calculus 3 to be correlated. All of the students completed calculus 2, but not all of the students completed calculus 3. To study the sensitivity of the random teacher effects from the GP VAM to the assumption that missing calculus 3 scores are MAR, Karl et al. (2013a) build a model for

student completion of calculus 3 (the "attendance model"). In their joint model referred to as MNAR-t, student attendance in calculus 3 is modeled as a function of effects associated with their calculus 2 teachers, recognizing that a teacher's effect on grades may be associated with the propensity of his or her future students to complete calculus 3. The attendance effect of each calculus 2 teacher is allowed to be correlated with the teacher's two effects from the test score model, and these effects are thus modeled with a 3×3 unstructured covariance matrix $\mathbf{\Gamma}^*_{teach,1}$ in the joint model. The student grades, $\mathbf{y}_{(1)}$, are assumed to be normally distributed and conditionally independent: the mixing of classmates from year to year leads to a multiple membership structure in $\mathbf{Z}_{(1)}$. The binary calculus 3 attendance indicators, $\mathbf{y}_{(2)}$, are modeled with a probit link. The $\mathbf{G}$ matrix for this example takes the form

$$\mathbf{G} = \text{cov}(\mathbf{b}) = \text{blockdiag}\left(\Gamma_{\text{stu}}, \ldots, \Gamma_{\text{stu}}, \mathbf{\Gamma}^*_{teach,1}, \ldots, \mathbf{\Gamma}^*_{teach,1}, \ldots, \Gamma_{teach,2} \ldots, \Gamma_{teach,2}\right)$$

We refer to Karl et al. (2013a) for a listing of parameter estimates and a discussion of the results. Table 3 lists the median iteration times for different approximations for this model.

TABLE 3. Median run times for iterations with first-order Laplace approximation, fully exponential corrections from (15) only, and fully exponential Laplace approximation.

Approximation	Iteration Time
First-order	1.5 min
FE:(15)-only	3.0 min
Fully Exponential (FE)	27 hr

6. SIMULATIONS

We illustrate the performance of our algorithm with two sets of simulations based on the binary model for the home team win/loss indicators. We focus on the binary model to highlight the performance gains available from the Laplace and fully exponential Laplace approximations in a non-nested application. Joe (2008) runs similar trials for nested binary and Poisson models. We have made both our simulation code and our code for the Poisson-binary model available in the online supplement, providing readers with the necessary tools to run more extensive simulations.

$$f(\mathbf{y}_{(2)}|\mathbf{b}) = \prod_{g=1}^{n} \left[\Phi\left\{(-1)^{1-y_{g(2)}}\left(\beta_{(2)} + \mathbf{z}'_{g(2)}\mathbf{b}\right)\right\}\right] \tag{19}$$

$$\mathbf{b} \sim N(\mathbf{0}, \sigma_w^2\mathbf{I}) \tag{20}$$

The only difference between this model and the one specified in (3) is the presence of the constant fixed effect, $\beta_{(2)}$. Since the $\mathbf{y}_{(2)}$ are indicators for a home-win, $\beta_{(2)}$ models the marginal probability of teams winning games played at home.

Each simulation consists of 500 runs with a vector of 100 "team ratings." In one simulation, the ratings are generated from a $N(0, 0.5)$ distribution, matching the structure assumed by the model. In the other simulation, the ratings are simulated from a t_3 distribution, and scaled by $1/\sqrt{6}$ to have variance approximately 0.5. Each "team" plays a random schedule of 4 home games and 4 away games. The binary win/loss outcomes are evaluated with model (19). A home field effect of 0.1 was built into the simulations, meaning, for example, the home team in a game between two evenly matched teams has a 55.6% chance of winning. We compare the performance of a pseudo-likelihood routine (using residual maximum likelihood), the first-order Laplace approximation, the first-order Laplace approximation with fully exponential corrections from Equation (15) only, and the fully exponential Laplace approximation. Our PL routine gives the same results that would be obtained by the default METHOD=RSPL setting of SAS PROC GLIMMIX. We have provided a copy of this program in the supplementary online material to accompany the Laplace routines.

The behavior of the parameter estimates in Tables 4 and 5 is consistent with the findings of previous work (Tierney et al., 1989; Breslow and Lin, 1995; Lin and Breslow, 1996; Steele, 1996). The RSPL parameter estimates are biased, especially the variance component estimate. The first-order and fully exponential Laplace approximations provide progressive improvements. The estimates using FE:(15)-only show much of the improvement of the fully exponential Laplace approximation in a fraction of the time.

TABLE 4. Median parameter estimates from 500 runs, using a $N(0, 0.5)$ distribution.

	Simulated Value	RSPL	Laplace	FE:(15)-only	FE Laplace
σ_w^2	0.5	0.309	0.335	0.479	0.506
$\beta_{(2)}$	0.1	0.084	0.092	0.098	0.100

TABLE 5. Median parameter estimates from 500 runs, using a $\frac{1}{\sqrt{6}}t_3$ distribution.

	Simulated Value	RSPL	Laplace	FE:(15)-only	FE Laplace
σ_w^2	0.5	0.207	0.224	0.295	0.310
$\beta_{(2)}$	0.1	0.088	0.094	0.097	0.098

TABLE 6. Median run times (seconds) for iterations with first-order Laplace approximation, fully exponential corrections from (15) only, and fully exponential Laplace approximation.

Approximation	Iteration Time
First-order	0.25
FE:(15)-only	0.42
Fully Exponential (FE)	1.45

7. SUMMARY

We demonstrate how the use of an EM algorithm with a fully exponential Laplace approximation proposed by Steele (1996) and Rizopoulos et al. (2009) may be adapted for the estimation of multiple response, non-nested, generalized linear mixed models. We illustrate the utility of the method in the context of a correlated random effects model proposed by Karl et al. (2013a) for joint modeling of student test scores and attendance, as well as with a Poisson-binary model for evaluating sports teams. We have provided sample code to build pseudo-likelihood, first-order Laplace, and fully exponential Laplace routines in the online supplementary material. Readers interested in modifying the R code for their own applications must primarily 1) construct the appropriate $\mathbf{Z}_{(i)}$ matrices for their model, 2) change the definitions of four functions which return the derivatives appearing in Section 3.1, and 3) modify the M step updates for the $\mathbf{G}$ matrix.

The presence of a linear predictor in GLMMs leads to a simplification that allows for the development of a template for joint estimation of an arbitrary number of GLMMs with different types of responses. In fact, there is no requirement here that the response distribution belong to the exponential family. As long as the mean of the distribution can be expressed as a function of a linear combination of the fixed and random effects, this template will work. The method could be extended to more general nonlinear mixed models; however, some simplifications, such as (7), would no longer apply in this setting.

The improved numerical accuracy from the algorithms presented in this paper are important for calculating the maximum likelihood estimates themselves, but come into play even more when the GLMM model-fitting is iterated in a jackknife or bootstrap. Lohr (2007) illustrates the sensitivity of bootstrap and jackknife calculations of small area estimates to the numerical method used to calculate parameter estimates of a GLMM. The methods in this paper, with their increased accuracy, are potentially useful in conjunction with resampling methods for variance estimation.

The proposed EM algorithm and computational routine may also be effectively applied to nested models. Aldor-Noiman et al. (2009) present a mixed model for forecasting the volume of calls received by a call center over the course of a day. The number of calls in each period is transformed to be approximately normal. Independently, the mean service time in each period is modeled using a normal distribution, and the predictions for call volume and service times are combined for an estimate of the workload faced by the center. It is feasible that the number of calls received in a period may be correlated with the length of those calls. The correlated random effects model described in this paper may be used to develop either a normal-exponential or a Poisson-exponential model for call center volume and service time.

Appendix A: Derivatives of the Conditional Log-Likelihoods

Poisson Response. For a Poisson response with a log link,

$$\begin{aligned}
\log f(\mathbf{y}_{(i)}|\mathbf{b}) &= \sum_{l=1}^{n_i} \left\{ -\log\left(y_{l(i)}!\right) + y_{l(i)}\eta_{l(i)} - \exp\left(\eta_{l(i)}\right) \right\} \\
\frac{\partial \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}} &= y_{l(i)} - \exp\left(\eta_{l(i)}\right) \\
\frac{\partial^2 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^2} &= \frac{\partial^3 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^3} = \frac{\partial^4 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^4} = -\exp\left(\eta_{l(i)}\right)
\end{aligned}$$

Binary Response. Suppose a binary response $\mathbf{y}_{(i)}$ is modeled with a probit link, and let $\Phi(\cdot)$ represent the standard normal distribution function.

$$\begin{aligned}
\log f(\mathbf{y}_{(i)}|\mathbf{b}) &= \sum_{l=1}^{n_i} \left\{ y_{l(i)} \log\left(\Phi\left[\eta_{l(i)}\right]\right) + \left(1 - y_{l(i)}\right) \log\left(1 - \Phi\left[\eta_{l(i)}\right]\right) \right\} \\
&= \sum_{l=1}^{n_i} \log\left(\Phi\left[(-1)^{1-y_{l(i)}}\, \eta_{l(i)}\right]\right)
\end{aligned}$$

Let $\eta^*_{l(i)} = (-1)^{1-y_{l(i)}}\, \eta_{l(i)}$ to compress notation,

$$\begin{aligned}
\frac{\partial \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}} &= (-1)^{1-y_{l(i)}}\, \frac{\phi(\eta^*_{l(i)})}{\Phi(\eta^*_{l(i)})} \\
\frac{\partial^2 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^2} &= \frac{\frac{\partial \phi(\eta^*_{l(i)})}{\partial \eta_{l(i)}} \Phi\left(\eta^*_{l(i)}\right) - \phi^2\left(\eta^*_{l(i)}\right)}{\Phi^2\left(\eta^*_{l(i)}\right)} \\
\frac{\partial^3 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^3} &= (-1)^{1-y_{l(i)}}\, \frac{\frac{\partial^2 \phi(\eta^*_{l(i)})}{\partial \eta^2_{l(i)}} \Phi^2(\eta^*_{l(i)}) - 3\frac{\partial \phi(\eta^*_{l(i)})}{\partial \eta_{l(i)}} \Phi(\eta^*_{l(i)})\phi(\eta^*_{l(i)}) + 2\phi^3\left(\eta^*_{l(i)}\right)}{\Phi^3\left(\eta^*_{l(i)}\right)} \\
\frac{\partial^4 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^4} &= \frac{\frac{\partial^3 \phi(\eta^*_{l(i)})}{\partial \eta^3_{l(i)}} \Phi^3\left(\eta^*_{l(i)}\right) - 4\frac{\partial^2 \phi(\eta^*_{l(i)})}{\partial \eta^2_{l(i)}} \Phi^2(\eta^*_{l(i)})\phi(\eta^*_{l(i)})}{\Phi^4\left(\eta^*_{l(i)}\right)} \\
&\quad + \frac{12\frac{\partial \phi(\eta^*_{l(i)})}{\partial \eta_{l(i)}} \Phi(\eta^*_{l(i)})\phi^2(\eta^*_{l(i)}) - 3\left(\frac{\partial \phi(\eta^*_{l(i)})}{\partial \eta_{l(i)}}\right)^2 \Phi^2(\eta^*_{l(i)}) - 6\phi^4(\eta^*_{l(i)})}{\Phi^4\left(\eta^*_{l(i)}\right)}
\end{aligned}$$

where $\phi(\cdot)$ is the standard normal density function and we use the notional convention that

$$\frac{\partial \phi(\eta^*_{l(i)})}{\partial \eta_{l(i)}} = \frac{\partial \phi(\eta_{l(i)})}{\partial \eta_{l(i)}}\Big|_{\eta_{l(i)} = \eta^*_{l(i)}}$$

Normal Response. For a normal response, $y_{l(i)} \sim N(\eta_{l(i)}, \sigma_i^2)$,

$$\begin{aligned}
\frac{\partial \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}} &= \frac{y_{l(i)} - \eta_{l(i)}}{\sigma_i^2} \\
\frac{\partial^2 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^2} &= -\frac{1}{\sigma_i^2} \\
\frac{\partial^3 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^3} &= \frac{\partial^4 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^4} = 0
\end{aligned}$$

Appendix B: Calculation of Terms for the E step

The E step calculations requires the terms $\partial\boldsymbol{\Sigma}^{(c)}/\partial c_k$, $\partial^2\boldsymbol{\Sigma}^{(c)}/\partial c_k\partial c_d$, and $\partial H(\mathbf{b})/\partial\mathbf{b}$. Furthermore, the calculation of the first two of these terms requires calculation of $\partial\widehat{\mathbf{b}}^{(c)}/\partial c_k$ and $\partial^2\widehat{\mathbf{b}}^{(c)}/\partial c_k\partial c_d$, both evaluated at $\boldsymbol{c}=0$. We first calculate $\partial\widehat{\mathbf{b}}^{(c)}/\partial c_k$, following the method of Rizopoulos et al. (2009). Let

$$\kappa(\mathbf{b}) = \left[\sum_{i=1}^{q}\log f\left(\mathbf{y}_{(i)}|\mathbf{b};\boldsymbol{\Psi}\right)\right] + \log\left\{f\left(\mathbf{b}\right)\right\}.$$

Since, by definition, $\widehat{\mathbf{b}}^{(c)} = \text{argmax}_{\mathbf{b}}[\log\left\{f(\mathbf{y}_{(1)},\ldots,\mathbf{y}_{(q)},\mathbf{b})\right\} + \boldsymbol{c}'H(\mathbf{b})]$, we have

$$\begin{aligned}\mathbf{0} &= \frac{\partial}{\partial\mathbf{b}}\left\{\kappa(\mathbf{b}) + \boldsymbol{c}'H(\mathbf{b})\right\}_{\mathbf{b}=\widehat{\mathbf{b}}^{(c)}}\\ &= \frac{\partial\kappa(\widehat{\mathbf{b}}^{(c)})}{\partial\mathbf{b}} + \frac{\partial\boldsymbol{c}'H(\widehat{\mathbf{b}}^{(c)})}{\partial\mathbf{b}}\end{aligned}$$

Taking the derivative with respect to c_k yields

$$\frac{\partial^2\kappa(\widehat{\mathbf{b}}^{(c)})}{\partial\mathbf{b}\partial\mathbf{b}'}\frac{\partial\widehat{\mathbf{b}}^{(c)}}{\partial c_k} + \frac{\partial\boldsymbol{e}_k'H(\widehat{\mathbf{b}}^{(c)})}{\partial\mathbf{b}} + \left(\frac{\partial}{\partial c_k}\left\{\frac{\partial H(\widehat{\mathbf{b}}^{(c)})}{\partial\mathbf{b}}\right\}\right)'\boldsymbol{c} = \mathbf{0}$$

Solving for $\partial\widehat{\mathbf{b}}^{(c)}/\partial c_k$ and evaluating at $c=0$ gives

$$\begin{aligned}\frac{\partial\widehat{\mathbf{b}}^{(c)}}{\partial c_k}|_{\boldsymbol{c}=0} &= \left(-\frac{\partial^2\kappa(\widehat{\mathbf{b}})}{\partial\mathbf{b}\partial\mathbf{b}'}\right)^{-1}\left(\frac{\partial\boldsymbol{e}_k'H\left(\widehat{\mathbf{b}}\right)}{\partial\mathbf{b}}\right)\\ &= \boldsymbol{\Sigma}^{-1}\left(\frac{\partial\boldsymbol{e}_k'H\left(\widehat{\mathbf{b}}\right)}{\partial\mathbf{b}}\right)\end{aligned}$$

We only need the terms $\partial^2\boldsymbol{\Sigma}^{(c)}/\partial c_k\partial c_d$ and $\partial^2\widehat{\mathbf{b}}^{(c)}/\partial c_k\partial c_d$ for the case $H(\mathbf{b})=\mathbf{b}$. These terms are used in the calculation of $\widetilde{\mathbf{v}}$. To find $\partial^2\widehat{\mathbf{b}}^{(c)}/\partial c_k\partial c_d$ where $H(\mathbf{b})=\mathbf{b}$, note

$$\begin{aligned}&\frac{\partial}{\partial c_d}\left\{\frac{\partial^2\kappa(\widehat{\mathbf{b}}^{(c)})}{\partial\mathbf{b}\partial\mathbf{b}'}\frac{\partial\widehat{\mathbf{b}}^{(c)}}{\partial c_k} + \frac{\partial\boldsymbol{e}_k'H(\widehat{\mathbf{b}}^{(c)})}{\partial\mathbf{b}} + \left(\frac{\partial}{\partial c_k}\left\{\frac{\partial H(\widehat{\mathbf{b}}^{(c)})}{\partial\mathbf{b}}\right\}\right)'\boldsymbol{c}\right\} = \mathbf{0}\\ &\Rightarrow \frac{\partial}{\partial c_d}\left\{\frac{\partial^2\kappa(\widehat{\mathbf{b}}^{(c)})}{\partial\mathbf{b}\partial\mathbf{b}'}\frac{\partial\widehat{\mathbf{b}}^{(c)}}{\partial c_k} + \boldsymbol{e}_k\right\} = \mathbf{0}\\ &\Rightarrow \frac{\partial^2\kappa(\widehat{\mathbf{b}}^{(c)})}{\partial\mathbf{b}\partial\mathbf{b}'}\frac{\partial^2\widehat{\mathbf{b}}^{(c)}}{\partial c_k\partial c_d} + \frac{\partial^3\kappa(\widehat{\mathbf{b}}^{(c)})}{\partial\mathbf{b}\partial\mathbf{b}'\partial\mathbf{b}'}\frac{\partial\widehat{\mathbf{b}}^{(c)}}{\partial c_d}\frac{\partial\widehat{\mathbf{b}}^{(c)}}{\partial c_k} = \mathbf{0}\\ &\overset{\boldsymbol{c}=\mathbf{0}}{\Rightarrow} \frac{\partial^2\widehat{\mathbf{b}}^{(c)}}{\partial c_k\partial c_d}|_{\boldsymbol{c}=\mathbf{0}} = \left(\frac{\partial^2\kappa(\widehat{\mathbf{b}})}{\partial\mathbf{b}\partial\mathbf{b}'}\right)^{-1}\frac{\partial\boldsymbol{\Sigma}}{\partial\mathbf{b}'}\frac{\partial\widehat{\mathbf{b}}}{\partial c_d}\frac{\partial\widehat{\mathbf{b}}}{\partial c_k}\\ &\Rightarrow \frac{\partial^2\widehat{\mathbf{b}}^{(c)}}{\partial c_k\partial c_d}|_{\boldsymbol{c}=\mathbf{0}} = \boldsymbol{\Sigma}^{-1}\left[\sum_{i=1}^{q}\sum_{l=1}^{n_i}\frac{\partial^3\log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial\eta_{l(i)}^3}\mathbf{z}_{l(i)}\mathbf{z}_{l(i)}'\left(\mathbf{z}_{l(i)}'\boldsymbol{\Sigma}^{-1}\boldsymbol{e}_d\right)\right]\boldsymbol{\Sigma}^{-1}\boldsymbol{e}_k.\end{aligned}$$

We may thus write

$$\frac{\partial\boldsymbol{\Sigma}^{(\boldsymbol{c})}}{\partial c_k}|_{\boldsymbol{c}=\mathbf{0}} = -\sum_{i=1}^{q}\sum_{l=1}^{n_i}\left[\frac{\partial^3\log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial\eta_{l(i)}^3}\left(\mathbf{z}_{l(i)}'\frac{\partial\widehat{\mathbf{b}}^{(c)}}{\partial c_k}|_{\boldsymbol{c}=0}\right)\mathbf{z}_{l(i)}\mathbf{z}_{l(i)}'\right] - \frac{\partial^2}{\partial\mathbf{b}\partial\mathbf{b}'}\left[\boldsymbol{e}_k'H(\widehat{\mathbf{b}})\right]$$

When $H(\mathbf{b})=\mathbf{b}$ and when $(\boldsymbol{c},\mathbf{b})$ are evaluated at $(\mathbf{0},\widehat{\mathbf{b}})$,

$$\begin{aligned}\frac{\partial^2 \boldsymbol{\Sigma}^{(\mathbf{c})}}{\partial c_k \partial c_d} =& \frac{\partial}{\partial c_d}\left[-\sum_{i=1}^{q}\sum_{l=1}^{n_i}\frac{\partial^3 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^3}\left(\mathbf{z}'_{l(i)}\frac{\partial \widehat{\mathbf{b}}^{(c)}}{\partial c_k}|_{\boldsymbol{c}=0}\right)\mathbf{z}_{l(i)}\mathbf{z}'_{l(i)}\right]_{\boldsymbol{c}=\mathbf{0}} \\ =& -\sum_{i=1}^{q}\sum_{l=1}^{n_i}\left\{\left[\mathbf{z}'_{l(i)}\frac{\partial^2 \widehat{\mathbf{b}}^{(c)}}{\partial c_k \partial c_d}|_{\mathbf{c}=0}\frac{\partial^3 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^3}\right.\right. \\ & \left.\left. + \mathbf{z}'_{l(i)}\left(\frac{\partial \widehat{\mathbf{b}}^{(c)}}{\partial c_k}|_{\boldsymbol{c}=0}\right)\mathbf{z}'_{l(i)}\left(\frac{\partial \widehat{\mathbf{b}}^{(c)}}{\partial c_d}|_{\boldsymbol{c}=0}\right)\frac{\partial^4 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^4}\right]\mathbf{z}_{l(i)}\mathbf{z}'_{l(i)}\right\}\end{aligned}$$

The fully exponential approximation for the integrals in the score function of $\boldsymbol{\beta}_{(i)}$ requires two additional terms: $\partial H(\mathbf{b})/\partial \mathbf{b}$ and $\partial^2 H(\mathbf{b})/\partial \mathbf{b}\partial \mathbf{b}'$. Note that these terms are not included in the summand over l.

$$\begin{aligned}\frac{\partial}{\partial \mathbf{b}}\left[\frac{\partial \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}}\right] &= \frac{\partial^2 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^2}\mathbf{z}_{l(i)} \\ \frac{\partial^2}{\partial \mathbf{b}\partial \mathbf{b}'}\left[\frac{\partial \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}}\right] &= \frac{\partial^3 \log[f(\mathbf{y}_{(i)}|\mathbf{b})]}{\partial \eta_{l(i)}^3}\mathbf{z}_{l(i)}\mathbf{z}'_{l(i)}\end{aligned}$$

Acknowledgments

This research was partially supported by the National Science Foundation under grant DRL-0909630, and by Arizona State University through a Dissertation Fellowship. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of the National Science Foundation, Adsurgo LLC, Arizona State University, or Westat. We would like to thank the anonymous reviewers, whose suggestions lead to substantial improvements in the paper.

References

Aldor-Noiman, S., Feigin, P. D., and Mandelbaum, A. (2009), "Workload Forecasting for a Call Center: Methodology and a Case Study," *The Annals of Applied Statistics*, 3, 1403–1447.

Annis, D. H. and Craig, B. A. (2005), "Hybrid Paired Comparison Analysis, with Applications to the Ranking of College Football Teams," *Journal of Quantitative Analysis in Sports*, 1, Iss. 1, Article 3.

Bates, D. and Maechler, M. (2012), *Matrix: Sparse and Dense Matrix Classes and Methods*, http://cran.r-project.org/web/packages/Matrix/index.html, R package version 1.0-4.

Bolker, B. M., Brooks, M. E., Clark, C. J., Geange, S. W., Poulsen, J. R., Stevens, M. H. H., and White, J.-S. S. (2009), "Generalized Linear Mixed Models: A Practical Guide for Ecology and Evolution," *Trends in Ecology and Evolution*, 24, 127–135.

Breslow, N. E. and Clayton, D. G. (1993), "Approximate Inference in Generalized Linear Mixed Models," *Journal of the American Statistical Association*, 88, 9–25.

Breslow, N. E. and Lin, X. (1995), "Bias Correction in Generalised Linear Mixed Models with a Single Component of Dispersion," *Biometrika*, 82, 81–91.

Broatch, J. and Lohr, S. (2012), "Multidimensional Assesment of Value Added by Teachers to Real-World Outcomes," *Journal of Educational and Behavioral Statistics*, 37, 256–277.

Browne, W. J., Goldstein, H., and Rasbash, J. (2001), "Multiple Membership Multiple Classification (MMMC) Models," *Statistical Modelling*, 1, 103–124.

Brumback, B. A., Zheng, H. W., and Dailey, A. B. (2013), "Adjusting for Confounding by Neighborhood Using Generalized Linear Mixed Models and Complex Survey Data," *Statistics in Medicine*, 32, 1313–1324.

Chandola, T., Clarke, P., Wiggins, R. D., and Bartley, M. (2005), "Who You Live With and Where You Live: Setting the Context for Health Using Multiple Membership Multilevel Models," *The Journal of Epidemiology and Community Health*, 59, 170–175.

de Leeuw, J. and Meijer, E. (2008), *The Handbook of Multilevel Analysis*, New York: Springer.

Demidenko, E. (2013), *Mixed Models: Theory and Applications with R*, Hoboken, NJ: Wiley, 2nd ed.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977), "Maximum Likelihood from Incomplete Data Via the EM Algorithm," *Journal of the Royal Statistical Society, Series B*, 39, 1–38.

Eugster, M. J. A., Knaus, J., Porzelius, C., Schmidberger, M., and Vicedo, E. (2011), "Hands-on Tutorial for Parallel Computing with R," *Computational Statistics*, 26, 219–239.

Evans, M. and Swartz, T. (1995), "Methods for Approximating Integrals in Statistics with Special Emphasis on Bayesian Integration Problems," *Statistical Science*, 10, 254–272.

Fieuws, S. and Verbeke, G. (2006), "Pairwise Fitting of Mixed Models for the Joint Modeling of Multivariate Longitudinal Profiles," *Biometrics*, 62, 424–431.

Gueorguieva, R. V. and Sanacora, G. (2006), "Joint Analysis of Repeatedly Observed Continuous and Ordinal Measures of Disease Severity," *Statistics in Medicine*, 25, 1307–1322.

Harville, D. A. (2008), *Matrix Algebra from a Statistician's Perspective*, New York: Springer.

Jamshidian, M. and Jennrich, R. I. (2000), "Standard Errors for EM Estimation," *Journal of the Royal Statistical Society, Series B*, 62, 257–270.

Joe, H. (2008), "Accuracy of Laplace Approximation for Discrete Response Mixed Models," *Computational Statistics and Data Analysis*, 52, 5066–5074.

Karl, A. T. (2012), "The Sensitivity of College Football Rankings to Several Modeling Choices," *Journal of Quantitative Analysis in Sports*, 8, Issue 3, doi:10.1515/1559–0410.1471.

Karl, A. T., Yang, Y., and Lohr, S. (2012), *GPvam: Maximum Likelihood Estimation of Multiple Membership Mixed Models Used in Value-Added Modeling*, http://cran.r-project.org/web/packages/GPvam/index.html, R package version 2.0-0.

Karl, A. T., Yang, Y., and Lohr, S. L. (2013a), "A Correlated Random Effects Model for Nonignorable Missing Data in Value-Added Assessment of Teacher Effects," *Journal of Educational and Behavioral Statistics*, 38, 577–603.

— (2013b), "Efficient Maximum Likelihood Estimation of Multiple Membership Linear Mixed Models, with an Application to Educational Value-Added Assessments," *Computational Statistics and Data Analysis*, 59, 13–27.

Kass, R. E. and Steffey, D. (1989), "Approximate Bayesian Inference in Conditionally Independent Hierarchical Models (Parametric Empirical Bayes Models)," *Journal of the American Statistical Association*, 84, 717–726.

Laird, N. M. and Ware, J. H. (1982), "Random-Effects Models for Longitudinal Data," *Biometrics*, 38, 963–974.

Lehmann, E. L. and Romano, J. (2010), *Testing Statistical Hypotheses*, New York: Springer, 3rd ed.

Lin, H., Liu, D., and Zhou, X.-H. (2009), "A Correlated Random-Effects Model for Normal Longitudinal Data with Nonignorable Missingness," *Statistics in Medicine*, 29, 236–247.

Lin, X. and Breslow, N. E. (1996), "Bias Correction in Generalized Linear Mixed Models With Multiple Components of Dispersion," *Journal of the American Statistical Association*, 91, 1007–1016.

Lindstrom, M. J. and Bates, D. M. (1988), "Newton-Raphson and EM Algorithms for Linear Mixed-Effects Models for Repeated-Measures Data," *Journal of the American Statistical Association*, 83, 1014–1022.

Little, R. and Rubin, D. (2002), *Statistical Analysis with Missing Data*, New York: John Wiley, 2nd ed.

Lohr, S. (2007), "Discussion of Papers on Resampling Methods in Small Area Estimation," in *Proceedings of the Survey Research Methods Section, American Statistical Association*, pp. 3265–3267.

Louis, T. A. (1982), "Finding the Observed Information Matrix when Using the EM Algorithm," *Journal of the Royal Statistical Society, Series B*, 44, 226–233.

Magnus, J. R. and Neudecker, H. (1999), *Matrix Differential Calculus with Applications in Statistics and Econometrics*, Hoboken, NJ: Wiley, revised ed.

Mariano, L. T., McCaffrey, D. F., and Lockwood, J. (2010), "A Model for Teacher Effects from Longitudinal Data Without Assuming Vertical Scaling," *Journal of Educational and Behavioral Statistics*, 35, 253–279.

McCullagh, P. and Nelder, J. A. (1989), *Generalized Linear Models*, London: Chapman & Hall, 2nd ed.

McCulloch, C. E. (1994), "Maximum Likelihood Variance Components Estimation for Binary Data," *Journal of the American Statistical Association*, 89, 330–335.

McLachlan, G. J. and Krishnan, T. (2008), *The EM Algorithm and Extensions*, Hoboken, NJ: Wiley, 2nd ed.

Page, K. and Page, L. (2010), "Alone Against the Crowd: Individual Differences in Referees Ability to Cope Under Pressure," *Journal of Economic Psychology*, 31, 192–199.

Pinheiro, J. C. and Bates, D. M. (1995), "Approximations to the Log-Likelihood Function in the Nonlinear Mixed-Effects Model," *Journal of Computational and Graphical Statistics*, 4, 12–35.

Pinheiro, J. C. and Chao, E. C. (2006), "Efficient Laplacian and Adaptive Gaussian Quadrature Algorithms for Multilevel Generalized Linear Mixed Models," *Journal of Computational and Graphical Statistics*, 15, 58–81.

R Core Team (2013), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.

Raudenbush, S. W., Yang, M.-L., and Yosef, M. (2000), "Maximum Likelihood for Generalized Linear Models with Nested Random Effects via High-Order, Multivariate Laplace Approximation," *Journal of Computational and Graphical Statistics*, 9, 141–157.

Rizopoulos, D. (2010), "JM: An R Package for the Joint Modelling of Longitudinal and Time-to-Event Data," *Journal of Statistical Software*, 35, 1–33.

Rizopoulos, D., Verbeke, G., and Lesaffre, E. (2009), "Fully Exponential Laplace Approximations for the Joint Modelling of Survival and Longitudinal Data," *Journal of the Royal Statistical Society, Series B*, 71, 637–654.

Shun, Z. (1997), "Another Look at the Salamander Mating Data: A Modified Laplace Approximation Approach," *Journal of the American Statistical Association*, 92, 341–349.

Shun, Z. and McCullagh, P. (1995), "Laplace Approximation of High Dimensional Integrals," *Journal of the Royal Statistical Society, Series B*, 57, 749–760.

Steele, B. M. (1996), "A Modified EM Algorithm for Estimation in Generalized Mixed Models," *Biometrics*, 52, 1295–1310.

Stroup, W. W. (2012), *Generalized Linear Mixed Models: Modern Concepts, Methods and Applications*, Boca Raton, FL: CRC Press.

Tierney, L. and Kadane, J. B. (1986), "Accurate Approximations for Posterior Moments and Marginal Densities," *Journal of the American Statistical Association*, 81, 82–86.

Tierney, L., Kass, R. E., and Kadane, J. B. (1989), "Fully Exponential Laplace Approximations to Expectations and Variances of Nonpositive Functions," *Journal of the American Statistical Association*, 84, 710–716.

Vonesh, E. F. (1996), "A Note on the Use of Laplace's Approximation for Nonlinear Mixed-Effects Models," *Biometrika*, 83, 447–452.

Wolfinger, R. and O'Connell, M. (1993), "Generalized Linear Mixed Models: A Pseduo-likelihood Approach," *Journal of Statistical Computation and Simulation*, 48, 233–243.

Wu, C.F.J. (1983), "On the Convergence Properties of the EM Algorithm," *Annals of Statistics*, 11, 95–103.

Wu, M. C. and Carroll, R. J. (1988), "Estimation and Comparison of Changes in the Presence of Informative Right Censoring by Modeling the Censoring Process," *Biometrics*, 44, 175–188.